

23. The apparatus of claim 21 wherein the second means computes a local address by adding an offset of one or more bits to the global address, yielding the local address.
24. The apparatus of claim 21 wherein the second means computes a local address by replacing one or more bits of the global address by a base value, yielding the local address.
25. The apparatus of claim 21 wherein the first means propagates the message to a number of processor elements, the number comprising less than all and more than one of the processor elements.

REMARKS

This paper responds to the Office Action dated March 12, 2001 (paper no. 14) in the parent application no. 09/121,493. To more particularly set forth the claimed invention, this Amendment cancels the previously pending claims and adds new claims. A rough correspondence of old and new claims is given in the following table:

old claim	1	3	5	6	8	13
new claim	16	17	18	19	20	21

The Examiner rejected all pending claims as supposedly obvious over Gifford U.S. Pat. No. 4,891,787 ("Gifford") if combined with Summer Jr. et al. U.S. Pat. No. 4,414,624 ("Summer").

Summer describes a *shared* memory architecture of a well-known kind. The Examiner is respectfully referred to the first two pages of the specification, in which the applicant acknowledges that systems with *shared* memories are well known. In a typical shared-memory system, the memories all reside on a parallel bus having address lines, data lines, and control lines. In such a system, any device on the bus (e.g. one of several processors) may read data from and write data to any other device on the bus (e.g. one of several memories) by means of a traditional bus read or bus write procedure. Such a bus has the advantage of having high bandwidth and low latency (see specification page 5, lines 20-22).

Such high bandwidth and low latency comes at a cost, or more accurately at several costs. There are limits on the practical length of such a bus, due to the non-infinite speed of propagation of signals on such a bus. Processor elements to be used on such a bus are expensive. Mechanisms must be provided to arbitrate in the event that more than one device wishes to use the bus at a particular time. Access to memory becomes a bottleneck and thus such systems are not easily scalable. A typical maximum system size is about 32 processors. (See specification page 2, lines 8-10.) They are, however, easy to program and are thus attractive to consider.

But the claimed system does not have *shared* memories of this type, across the various processors. As may be seen in claims 16 and 21, there is no common bus connecting the various

memories. Instead, message-passing is employed to pass data between and among the processors.

Message passing, in and of itself, is of course also well known and is well suited to many applications. It is commonplace to encounter systems having as many as ten thousand processors (see specification page 2, lines 14-15). But historically, message passing has not proven completely well suited to use in parallel processing systems such as those discussed herein. For example, when commonly available operating systems are employed in prior-art message-passing systems, the latency is poor. It is not easy to program such systems. (Specification page 2, lines 15-20.)

The (amended) claims set forth arrangements that are different from either of the arrangements discussed above and in the cited prior art reference Summer. The claimed invention does not rely on there being a common bus connecting the memories of the various processor elements, but may equally well be used in networks without a common bus. It thus uses message-passing to pass computational results to and from the memories, yet *avoids* the latency problems of prior-art message-passing systems. The claims are expressed, however, in structural terms, not in terms of particular results. It is submitted that the structural limitations in the independent claims are not rendered unpatentable by the cited references.

Returning to the Summer reference, it uses a common bus among the memory elements of the various processor elements. See for example Fig. 1 "shared memory bus" 26. As such it has little or nothing in common with the Gifford system, nor with the claimed arrangement which lacks such a common bus. It is respectfully suggested that one skilled in the art would not try to integrate into a single system two systems as unlike as "oil and water," namely to try to integrate a shared-memory system into a message-passing system. Such systems have different physical arrangements and use different programming techniques.

It should be noted that the method according to the invention brings about unobvious advantages. For example, there are no data consistency problems, and fewer synchronization problems than in shared systems according to the state of the art. Since there is no reading path, there are no data consistency keeping problems, no need for arbitration for reading, thus no need for a reading arbiter nor any other reading controls.

By reason of the foregoing, it is requested that the amended claims be considered and allowed.

Respectfully submitted,



Carl Oppedahl
PTO Reg. No. 32,746
Oppedahl & Larson LLP
P O Box 5068
Dillon, CO 80435-5068
email oppedahl@patents.com